

1 METHODS AND APPARATUS FOR TRANSPORTING A SYNCHRONOUS OR
2 PLESIOCHRONOUS SIGNAL OVER A PACKET NETWORK
3

4 This application claims the benefit of co-owned, co-pending,
5 provisional application serial number 60/276,630 filed March 16,
6 2001, entitled "Methods and Apparatus for Delivering Multimedia
7 Communications Services to Multiple Tenants in a Building," the
8 complete disclosure of which is hereby incorporated by reference
9 herein.

10
11 BACKGROUND OF THE INVENTION
12

13 1.. Field of the Invention

14 The invention relates to telecommunications. More
15 particularly, the invention relates to clock recovery when
16 transporting a synchronous or plesiochronous signal over a packet
17 network.
18

19 2. State of the Art

20 The first commercial digital voice communications system was
21 installed in 1962 in Chicago, Illinois. The system was called
22 "T1" and was based on the time division multiplexing (TDM) of
23 twenty-four telephone calls on two twisted wire pairs. The
24 digital bit rate of the T1 system was 1.544 Mbit/sec (± 200 bps),
25 which was, in the nineteen sixties, about the highest data rate

1 that could be supported by a twisted wire pair for a distance of
2 approximately one mile. The cables carrying the T1 signals were
3 buried underground and were accessible via manholes, which were,
4 at that time in Chicago, spaced approximately one mile apart.
5 Thus, analog amplifiers with digital repeaters were conveniently
6 located at intervals of approximately one mile.

7
8 The T1 system is still widely used today and forms a basic
9 building block for higher capacity communication systems including
10 T3 which transports twenty-eight T1 signals. The designation T1
11 was originally coined to describe a particular type of carrier
12 equipment. Today T1 is often used to refer to a carrier system, a
13 data rate, and various multiplexing and framing conventions.
14 While it is more accurate to use the designation "DS1" when
15 referring to the multiplexed digital signal carried by the T1
16 carrier, the designations DS1 and T1 are often used
17 interchangeably. Today, T1/DS1 systems still have a data rate of
18 1.544 Mbit/sec and support typically twenty-four voice and/or data
19 DS0 channels. Similarly, the designations DS2 and T2 both refer
20 to a system transporting up to four DS1 signals (96 DS0 channels)
21 and the designations DS3 and T3 both refer to a system
22 transporting up to seven DS2 signals (672 DS0 channels). The
23 timing tolerance for modern T1 equipment has been raised to ± 50
24 bps. T1 signals are said to be "plesiochronous" (nearly
25 synchronous). Clock variations at nodes are corrected by line

1 codes such as alternate mark inversion (AMI or bipolar line code).
2 These codes set up patterns in the bitstream of the signal which
3 are used at nodes to correct for clock variations.
4

5 Today, another higher bandwidth TDM system is in use. This
6 system is referred to as the synchronous optical network (SONET)
7 or, in Europe, the synchronous digital hierarchy (SDH). Unlike
8 plesiochronous signals, SONET signals are synchronized to a master
9 network clock. Although the timing of SONET signals is very
10 accurate, some clock variations still exist at different nodes in
11 the network. Various complex techniques are provided to correct
12 for clock differences at different nodes.
13

14 The T1 and T3 networks were originally designed for digital
15 voice communication. In a voice network minor bit errors can be
16 tolerated as a small amount of noise. However, in a data network,
17 a minor bit error cannot be tolerated. In the early 1970s,
18 another technology was deployed to support data networks. The
19 technology was called "packet switching". Unlike the T1 and T3
20 networks, packet switching was designed for data communications
21 only. In packet switching, a "packet" of data includes a header,
22 a payload, and a cyclic redundancy check (CRC). The header
23 includes addressing information as well as an indication of the
24 length of the payload. The payload contains the actual data which
25 is being transmitted over the network. The CRC is used for error

1 detection. The receiver of the packet performs a calculation with
2 the bits in the packet and compares the result of the calculation
3 to the CRC value. If the CRC value is not the same as the result
4 of the calculation, it means that the packet was damaged in
5 transit. According to the packet switching scheme, the damaged
6 packet is discarded and the receiver sends a message to the
7 transmitter to resend the packet. One popular packet switching
8 scheme for wide area networks (WANs), known as X.25, utilizes a
9 packet which has a fixed payload of 128 octets. Other packet
10 switching schemes allow variable length packets up to 2,000
11 octets. Frame Relay is an example of a WAN packet switching
12 scheme which utilizes variable sized packets and Ethernet is an
13 example of a local area network (LAN) packet switching scheme
14 which utilizes variable sized packets. Packet switching networks
15 are asynchronous and, by design, are not well suited for the
16 transmission of a streaming signal such as voice or video. If
17 streaming voice or video is transmitted via packets, small amounts
18 of noise in the signal will result in discontinuity of the stream,
19 echo, and other problems.

20
21 Concurrent with the development of packet switching several
22 groups around the world began to consider standards for the
23 interconnection of computer networks and coined the term
24 "internetworking". The leading pioneers in internetworking were
25 the founders of ARPANET (the Advanced Research Projects Network).

1 ARPA, a U.S. Department of Defense organization, developed and
2 implemented the transmission control protocol (TCP) and the
3 internet protocol (IP). The TCP/IP code was dedicated to the
4 public domain and was rapidly adopted by universities, private
5 companies, and research centers around the world. An important
6 feature of IP is that it allows fragmentation operations, i.e. the
7 segmentation of packets into smaller units. This is essential to
8 allow networks which utilize large packets to be coupled to
9 networks which utilize smaller packets. Today, TCP/IP is the
10 foundation of the Internet. It is used for email, file transfer,
11 and for browsing the Worldwide Web. It is so popular that many
12 organizations are hoping to make it the worldwide network for all
13 types of communication, including voice and video.

14
15 Perhaps the most awaited, and now fastest growing technology
16 in the field of telecommunications is known as Asynchronous
17 Transfer Mode (ATM) technology. ATM was originally conceived as a
18 carrier of integrated traffic, e.g. voice, data, and video. ATM
19 utilizes fixed length packets (called "cells") of 53 octets (5
20 octets header and 48 octets payload). ATM may be implemented in
21 either a LAN or a WAN. For ideal data transfer, ATM is used end
22 to end from the data source to the data receiver.

23
24 Current ATM service is offered in different categories
25 according to a user's needs. Some of these categories include

1 constant bit rate (CBR), variable bit rate (VBR), unspecified bit
2 rate (UBR), and available bit rate (ABR). CBR service is given a
3 high priority and is used for streaming data such as voice and
4 video where a loss of cells would cause a noticeable degradation
5 of the stream. UBR and ABR services are given a low priority and
6 are used for data transfers such as email, file transfer, and web
7 browsing where sudden loss of bandwidth (bursty bandwidth) can be
8 tolerated. ATM service is sometimes referred to as "statistical
9 multiplexing" as it attempts to free up bandwidth which is not
10 needed by an idle connection for use by another connection.

11
12 ATM switches (like other packet switches) typically include
13 multiple buffers, queues, or FIFOs for managing the flow of ATM
14 cells through the switch. Generally, a separate buffer is
15 provided for each outlet from the switch. However, it is also
16 known to have separate buffers at the inlets to the switch.
17 Buffer thresholds are set to prevent buffer overflow. If the
18 number of cells in a buffer exceeds the threshold, no more cells
19 are allowed to enter the buffer. Cells attempting to enter a
20 buffer which has reached its threshold will be discarded.

21
22 Whenever a synchronous or plesiochronous signal is
23 transmitted over a packet network, e.g. ATM or the Internet, the
24 originating clock must be recovered at the receiver. Clock
25 recovery is necessary to prevent excessive packet loss, to prevent

1 unacceptable delay in processing the signal, and, in the case of
2 TDM signals, to facilitate framing. One method of recovering a
3 clock signal, called adaptive clock recovery, involves measuring
4 the depth of a (jitter) buffer. If the buffer depth is greater
5 than a maximum threshold or is increasing with time, the local
6 clock rate is increased to cause the buffer to drain more quickly.
7 If the buffer depth is less than a minimum threshold or is
8 decreasing with time, the local clock rate is decreased to cause
9 the buffer to drain less quickly. The main drawback of this clock
10 recovery method is that following an adjustment of the clock rate,
11 there is a delay before the buffer depth stabilizes due to the
12 inertia of the buffer depth. This delay may cause instability or
13 excessive jitter in the recovered clock.

14
15 In January 1997, the ATM forum defined "Circuit Emulation
16 Service" (CES) as af-vtoa-0078.000. CES uses ATM AAL1 adaptation
17 to segment incoming E1 or T1 traffic into ATM cells with the
18 necessary timing information to ensure that the circuit can be
19 correctly reassembled at the destination. The timing information
20 is provided in the ATM cell headers and is referred to as the
21 synchronous residual time stamp (SRTS). The time stamp is used by
22 the receiver to determine the difference between a common
23 reference clock and the sender's local clock. SRTS assumes the
24 availability of a common synchronous network clock from which the
25 sender and receiver can both reference. It also assumes that the

1 T1/E1 signal enters an ATM network and remains in the ATM network
2 until it exits as a T1/E1 signal. If the signal passes through
3 other networks (e.g. IP networks or Ethernet networks) and loses
4 traceability to the common reference clock, SRTS fails. For
5 example, the previously incorporated co-owned application
6 describes a system in which ATM cells containing packetized T1/E1
7 signals are transported over Ethernet. The Ethernet receiver
8 cannot utilize SRTS to recover the clock of T1/E1 signals.
9

10 SUMMARY OF THE INVENTION

11
12 It is therefore an object of the invention to provide methods
13 and apparatus for transporting a synchronous or plesiochronous
14 signal over a packet network.
15

16 It is also an object of the invention to provide methods and
17 apparatus which do not rely on a common synchronous network clock
18 for transporting a synchronous or plesiochronous signal over a
19 packet network.
20

21 It is another object of the invention to provide methods for
22 transporting a synchronous or plesiochronous signal over a packet
23 network which do not suffer the disadvantages of prior art
24 adaptive clock recovery methods.
25

1 It is still another object of the invention to provide
2 methods and apparatus for transporting a synchronous or
3 plesiochronous signal over a packet network which operates over
4 any type of packet network including a hybrid network which
5 combines different types of packet switching between source and
6 destination.

7
8 In accord with these objects which will be discussed in
9 detail below, the methods and apparatus of the present invention
10 are exemplified with reference to a network in which a T1 signal
11 has been segmented into ATM AAL1 cells. The methods of the
12 invention include providing incoming and outgoing cell counters at
13 the "local" user-network-interface (UNI) where the AAL1 cells are
14 to be reassembled into a T1 signal. The invention is implemented
15 under two assumptions. The first assumption is that the "remote"
16 network-network-interface (NNI) receives and transmits AAL1
17 packets at a consistent rate. The second assumption is that when
18 the local UNI clock and remote NNI clock are locked, the number of
19 cells received at the UNI should increase at the same rate as the
20 number of cells transmitted by the UNI. According to the basic
21 method of the invention, the UNI is provided with an adjustable
22 clock and the clock rate is adjusted by comparing the incoming
23 cell count with the outgoing cell count. In particular, if the
24 outgoing cell count is smaller than the incoming cell count, the
25 clock rate is increased. If the outgoing cell count is larger

1 than the incoming cell count, the clock rate is decreased. In
2 order to minimize delay in clock adjustments, a "gear shift"
3 adjustment algorithm is provided. In the presently preferred
4 embodiment, four different levels of adjustment are provided,
5 level 3 being the coarsest (fastest) and level 0 being the finest
6 (slowest).

7
8 The apparatus of the invention includes a phase locked loop
9 (PLL) embodied in a programmable logic device (PLD). The
10 apparatus has run-time clock adjustment capabilities. According
11 to the presently preferred embodiment, the CPU performing AAL1
12 processing dynamically adjusts its own clock by reading a register
13 in the PLD. The register is referred to as CLKADJ and the value
14 written to it by the PLL is an absolute number of clock ticks to
15 add or subtract per million. The cell counters are preferably
16 implemented as 16-bit unsigned integer counters. The transmit
17 cell counter is incremented whenever a cell is sent to the NNI
18 from the UNI and the receive cell counter is incremented whenever
19 a cell is added to the receive buffer in the UNI. The cell
20 counters are modulo-65536 counters which wrap to zero and continue
21 to count up from 1.

22
23 To allow for network jitter and the relatively slow drift
24 rate of the cell counters in the presently preferred
25 implementation, the PLL routine is run at relatively long

1 intervals, e.g 2048 ms. The interval can be adjusted for
2 different applications. The PLL routine begins by computing the
3 difference between the receive cell counter and the transmit cell
4 counter. According to one embodiment, the difference is written
5 to the register CLKADJ. Thus, if, e.g., the UNI has received one
6 more cell than it has transmitted, the value 1 will be written to
7 CLKADJ. This will cause the local clock rate to be increased by
8 one tick per million. As mentioned above, in order to decrease
9 the convergence time of the PLL, the preferred embodiment of the
10 invention utilizes a "gear shift" algorithm. According to the
11 presently preferred embodiment, the value in the register CLKADJ
12 is magnified by a factor of 2^{Level} where $0 \leq Level \leq 3$. When $Level=3$, the
13 PLL will approach the correct frequency very quickly and will
14 eventually begin to "circle" the true frequency. According to the
15 "gear shift" algorithm of the invention, when circling is
16 detected, $Level$ is lowered to 2, then 1 and then 0. According to
17 the presently preferred embodiment, circling is detected by taking
18 the derivative of the CLKADJ value and detecting sign changes in
19 the derivative. After the derivative changes sign several times,
20 circling is detected and the level is decreased. When the level
21 reaches 0 it is no longer changed.

22
23 Additional objects and advantages of the invention will
24 become apparent to those skilled in the art upon reference to the

1 detailed description taken in conjunction with the provided
2 figures.

3 4 BRIEF DESCRIPTION OF THE DRAWINGS

5
6 Figure 1 is a simplified block diagram of a system
7 incorporating the invention;

8
9 Figure 2 is a simplified block diagram of portions of the UNI
10 according to the invention;

11
12 Figure 3 is a flow chart illustrating a first method of the
13 invention; and

14
15 Figure 4 is a flow chart illustrating a second method of the
16 invention.

17 18 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

19
20 Referring now to Figure 1, an apparatus incorporating the
21 invention generally includes a network-network-interface NNI 10
22 which is coupled to the PSTN 12 and to an Ethernet LAN 14, and at
23 least one user-network interface UNI 16 which is coupled to the
24 LAN 14 and a plurality of user devices such as telephones 18a,
25 18b, ...18n. The LAN 14 is a packet network and is substantially

1 the same as the LAN described in the previously incorporated co-
2 owned, co-pending provisional application wherein ATM cells
3 containing packetized T1/E1 signals are transmitted over an
4 Ethernet LAN.

5
6 According to the invention, the UNI 16 includes an AAL1
7 processor as well as an apparatus for recovering the clock rate of
8 a synchronous or plesiochronous signal carried from the PSTN over
9 the packet LAN 14 via the NNI 10. Turning now to Figure 2, the
10 AAL1 processor 20 includes a clock generator 22 and a CPU 24. The
11 CPU 24 performs the AAL1 processing and controls the clock
12 generator 22 by adding or removing a number of clock ticks per
13 million. The UNI is further provided with a programmable logic
14 device (PLD) 26 for implementing a phase locked loop (PLL) to
15 provide the CPU 24 with a clock tick number to adjust the CLOCK
16 22. The PLD 26 includes a number of counters, registers,
17 arithmetic processes and logical processes. More particularly,
18 the PLD 26 includes a received cell counter (CCrx) 28, a
19 transmitted cell counter (CCTx) 30, a subtractor 32 for
20 subtracting the number of transmitted cells from the number of
21 received cells (CCrx-CCTx), and a multiplier 34 for multiplying
22 the difference (CCrx-CCTx) times 2^{Level} , where "Level" is an integer
23 from zero to three.

1 The result of the computation $(CCrx - CCtx) \times 2^{Level}$ is stored in a
2 register 36 called CLKADJ. The number written in CLKADJ is an
3 absolute number of clock ticks per million for the CPU 24 to add
4 to the clock generator 22. The cell counters 28, 30 are
5 preferably implemented as 16-bit unsigned counters. The transmit
6 cell counter CCtx is incremented whenever a cell is sent to the
7 NNI (10 in Figure 1) from the UNI (16 in Figure 1) and the receive
8 cell counter CCrX is incremented whenever a cell is added to the
9 receive buffer (not shown) in the UNI. The cell counters are
10 modulo-65536 counters which wrap to zero and continue to count up
11 from 1. To allow for network jitter and the relatively slow drift
12 rate of the cell counters in the presently preferred
13 implementation, the calculation of CLKADJ is performed every 2048
14 ms. The interval can be adjusted for different applications. The
15 method and apparatus described thus far provides a simple PLL
16 based on the relative number of transmitted and received cells.
17 As mentioned above, however, this simple PLL may take a long time
18 lock.

19
20 According to the preferred embodiment of the invention, a
21 "gear shift" routine is also run each time CLKADJ is recomputed.
22 The remainder of the blocks in Figure 2 are used to implement the
23 gear shift routine which shifts the value of Level from an initial
24 value of three to a locked value of zero. These additional
25 components include a register 38 for storing the last previous

1 value "PrvCLKADJ" of CLKADJ, a subtractor 40 for determining the
2 change "ClkDiff" in CLKADJ by calculating (CLKADJ-PrvCLKADJ), a
3 register 42 for storing the calculated ClkDiff, and a register 44
4 for storing the previous value "PrvClkDiff" of ClkDiff. A
5 comparator 46 is provided for comparing ClkDiff and PrvClkDiff.
6 Based on the comparison, logic 48 may increment a circle counter
7 50 and/or trigger the reset 52 of PrvCLKADJ and PrvClkDiff. Logic
8 54 monitors the circle counter 50. When the circlecount has
9 reached a threshold, logic 54 triggers the reset 56, level
10 decrementer 58, and cell count doubler 60. The operation of the
11 components in Figure 2 are better understood by reference to the
12 complete PLL and gear shift routine which is illustrated in
13 Figures 3 and 4.

15 Turning now to Figure 3, the operation of the invention
16 begins at 100 when a connection is established. The exponent
17 "Level" is set to "3" at 102. The cell counters are read at 104
18 and CLKADJ and ClkDiff are calculated at 106 and 108 respectively.
19 If it is determined at 110 that the ClkDiff is zero, the routine
20 returns to check the value of Level at 118, read the cell counters
21 at 104, and recalculate CLKADJ and ClkDiff. As mentioned above,
22 the calculation of CLKADJ and ClkDiff is performed approximately
23 every two seconds, e.g 2048 ms, thus, there will be a delay
24 between 118 and 104 in the flow chart of Figure 3. It will be
25 appreciated that the first iteration of this method will result in

1 no ClkDiff which will be treated as ClkDiff=0. If the ClkDiff is
2 not zero, the most significant bit (MSB), i.e. the sign \pm , of
3 ClkDiff is compared to the MSB of PrvClkDiff at 112. If the signs
4 are different it means that the CLKADJ is "circling" the correct
5 adjustment value, i.e. moving up and down about the correct value.
6 In this case, the circle counter is incremented at 114. In either
7 case, the PrvClkDiff and PrvCLKADJ are reset to ClkDiff and
8 CLKADJ, respectively at 116.

9

10 The components for performing the method of Figure 3 are
11 shown in Figure 2 as the logic 48, the CircleCount register 50,
12 and the reset block 52. Before resuming the next iteration of the
13 method of Figure 3 at 104, the value of Level is checked at 118.
14 The "Level Check" algorithm also includes a Level adjustment
15 method which is illustrated in Figure 4.

16

17 Entering Figure 4 from block 118 in Figure 3, the Level Check
18 starts at 200. If it is determined at 202 that Level=0, no
19 further processing is needed and the routine returns at 204 to
20 block 104 in Figure 3. If Level \neq 0, it is determined at 206
21 whether the circlecount \geq 4. If it is not, the routine returns at
22 204. If it is, a number of operations are performed at 208 before
23 returning at 204 to block 104 in Figure 3. The first operation is
24 that the value of Level is decremented by one. It will be
25 appreciated that this is performed by the component 58 in Figure

1 2. The next operation is that the values of the cell counters are
 2 doubled. This is performed by the component 60 in Figure 2. This
 3 is done only once at each level decrement to smooth the transition
 4 from one level to the next. The last operation is the resetting
 5 of the CircleCount to zero which is performed by the component 56
 6 in Figure 2.

7

8 Exemplary pseudocode for performing the method of Figure 3 is
 9 listed below.

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

```

CLKADJ=(CCrx-CCtx<<Level;
If (CLKADJ-PrvCLKADJ!=0)
    Then
        ClkDiff=CLKADJ-PrvCLKADJ;
        If ((ClkDiff^PrvClkDiff)&0x8000)
            Then
                CircleCount++;
            End If
        PrvClkDiff=ClkDiff;
        PrvCLKADJ=CLKADJ;
    End If

```

23 Exemplary pseudocode for performing the method of Figure 4 is
 24 listed below.

25

26

27

28

```

If (CircleCount>=4&&Level>0
    Then
        Level=Level-1;

```

```
1          CCrx=CCrx<<1;
2          CCTx=CCTx<<1;
3          CircleCount=0;
4      EndIf
5
```

6 The methods and apparatus of the invention include providing
7 incoming and outgoing cell counters at the "local" user-network-
8 interface (UNI) and using these counters to reconstruct the clock
9 at the network-network-interface (NNI). According to the basic
10 method of the invention, the UNI is provided with an adjustable
11 clock and the clock rate is adjusted by comparing the incoming
12 cell count with the outgoing cell count. In particular, if the
13 outgoing cell count is smaller than the incoming cell count, the
14 clock rate is increased. If the outgoing cell count is larger
15 than the incoming cell count, the clock rate is decreased. In
16 order to minimize delay in clock adjustments, a "gear shift"
17 adjustment algorithm is provided. In the presently preferred
18 embodiment, four different levels of adjustment are provided,
19 level 3 being the coarsest (fastest) and level 0 being the finest
20 (slowest).

21
22 There have been described and illustrated herein methods and
23 apparatus for transporting a synchronous or plesiochronous signal
24 over a packet network. While particular embodiments of the
25 invention have been described, it is not intended that the
26 invention be limited thereto, as it is intended that the invention

1 be as broad in scope as the art will allow and that the
2 specification be read likewise. Thus, while particular hardware
3 and software have been disclosed, it will be appreciated that
4 other hardware and/or software could be utilized. Also, while
5 certain level values and circle count values have been shown, it
6 will be recognized that other values could be used with similar
7 results obtained depending on the application. Moreover, while
8 particular examples have been disclosed in reference to ATM,
9 Ethernet, and T1/E1 signals, it will be appreciated that the
10 invention is applicable to any packet switching network carrying
11 packetized synchronous or plesiochronous signals. Thus, while the
12 disclosure refers to AAL1 and "cell counters", when applied to
13 another packet network, "packet counters" will be used and another
14 packetizing scheme other than AAL1 will be used. It will
15 therefore be appreciated by those skilled in the art that yet
16 other modifications could be made to the provided invention
17 without deviating from its spirit and scope as so claimed.